

USER SUPPORT FOR AN MPP, JPL BEGINS THE PROCESS

Barbara Horner-Miller

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

At Supercomputing 92, it was announced that JPL/Caltech had entered into a collaborative agreement with Cray Research to become one of a few "Cray Centers of Excellence in Parallel Computing". Under this agreement a T3D will be installed at JPL in the fall of 1993. This computer will have 256 processors, 512 MW of memory and will be capable of peak speeds of 38 gigaflops. Preliminary user services planning for the new system is beginning. This presentation will go through some general questions expected to help us determine how we can best serve users on the new system. It will give our answers to those questions and indicate what this means to our planning process.

Background

Since 1989, the Jet Propulsion Laboratory (JPL) has gone through several hardware changes with a major change anticipated in October 1993. JPL's first supercomputer, a CRAY X-MP/18, was installed at the Information Processing Center at the Jet Propulsion Laboratory in Pasadena California in July 1989. This system was leased from Cray Research, Inc., by the California Institute of Technology (Caltech) to provide a high-performance supercomputing resource to scientists and engineers at JPL and the Caltech campus and to their collaborators for solving computationally intensive problems. This system supplemented an extensive network of computers already in place at JPL and the Caltech campus.

In July 1991, the X-MP/18 was replaced by a CRAY Y-MP2E/116 which will be upgraded to a Y-MP2E/232 in early April 1993. In October 1993, the major step of acquiring a T3D will be taken. We expect to remove the CRAY Y-MP2E in April 1994,

leaving the T3D as the production supercomputer for the JPL/Caltech community.

The Caltech/JPL user community represents the wide range of disciplines one would expect to find at a facility composed of a national laboratory and a prominent academic institution. The majority of the JPL users are scientists and engineers, while the majority of the Caltech campus users are researchers in physics and chemistry. The multi-disciplinary environment includes astrophysics, space physics, earth science, planetary exploration and science, telecommunications, and microgravity.

Questions Affecting Implementation

Once the need for a major upgrade or change is recognized, but before jumping into the actual process, one must carefully consider the clientele, purpose, and expected mode of operation which will result. These variables directly impact the type of user support, additional equipment and software needed. Some of the questions that should be addressed are presented below.

Who Will the Users Be?

The targeted users must be carefully considered. The users (and the funding) are the basis for all implementation decisions. Key issues for determining how much and what type of support users will require include the diversity and sophistication of the user base and the expectations of the users.

Do the targeted users represent a diverse mix of disciplines, a few disciplines or only a single discipline? The more similar the applications represented by the user base, the easier to establish the training program, design specific tools, and select suitable software. If the user base is reasonably homogeneous, users can derive additional support from each other as they learn. Having a limited number of disciplines to support also makes it easier to develop local tools which can service large portions of the user community and to structure training classes which stress the necessary skills.

Which specific disciplines will be supported under the new facility? Do special tools or software products already exist to support those disciplines? For example, many users work in quantum chemistry and CRI has announced that to be one of their targeted applications priorities. This implies the availability of additional tools in that area.

In addition to the scientific disciplines practiced by the users, there are other aspects of their similarity to be considered. Will the users be scientists and engineers or will they be professional programmers working with scientists and engineers? This affects the amount of time the average user can be expected to devote to learning the available hardware and software tools. Professional programmers are more likely to spend time learning the intricacies of software and the operating system than are engineers for whom the computer is primarily a research tool.

Are there other abilities that should be considered? Have the users had experience with any type of parallel processing? If so, is that experience with the use of two processors on a CRAY, or is it experience with other massively parallel systems? Answers to these questions will provide valuable assistance for gearing documentation and training efforts to the knowledge base already held by the users.

At JPL, the user base will be composed of users with diverse scientific disciplines. Most of these users will be scientists and engineers as opposed to professional programmers. We can expect them to be divided into two camps based on previous parallel processing experience.

The first camp consists of participants in the Center of Excellence agreement. These users generally have lots of experience with other massively parallel processors and are very sophisticated users. They are anxiously awaiting the arrival of the T3D to see how it compares with other MPPs they have used such as the Mark III^{HP} Hypercube, the CM-2, the CM-5, Intel Gamma, the Intel Touchstone Delta and Intel Paragon systems. Most of these users have codes that work on several MPP platforms. They have made explicit commitments to port and extend these codes for the CRAY T3D. The user services challenge facing us with this group of users is that they are considerably more sophisticated in this area than the user services staff is.

The second camp consists of users with little or no experience in parallel processing. These users are also anticipating the arrival of the CRAY T3D but they will be the test of CRI's desire to provide a user friendly MPP system. Those users need tutorial level documentation and basic training in parallelizing codes. Many will need hand holding as they port their first code to an MPP platform. The user services challenges facing us with this group of users are more familiar: providing a complete training course to introduce users to the tools and concepts needed for MPP programming; providing basic documentation; and providing assistance to users as they port their codes. These challenges although familiar will be difficult as we struggle to stay ahead of the users.

What is the Available Level of Funding?

The available level of funding obviously impacts the decisions on the hardware, software and staffing. Unless there are truly unlimited resources (and there very seldom, if ever, are), tradeoffs must be made between competing uses for funds.

The less sophisticated the programming knowledge of the users, the more user support required. The more sophisticated the hardware and software, the more user support required, regardless of the level

of user programming knowledge. When users are expected to do everything themselves, the hardware and the software must be relatively simple and the **available** resources must be channeled towards training and programming advice for the user. When, on the other hand, users are expected to provide mainly the knowledge of their program and the scientific interactions taking place, then available resources must be applied to the staff who will perform the actual tasks. In the later case, the software and hardware may be much more **sophisticated** because they will be used primarily by a small, well-trained staff. There are many levels of support between these extremes and the correct level changes over time as users adjust to the system.

Funding can be used to simplify the process for the user, but it can also have the opposite effect, that of making the process more difficult for the user, when is not applied wisely. Simplification might be achieved by purchasing third party software to provide additional MPP programming models; training tools; etc. Permanent or temporary staff with extensive MPP experience might be hired to assist users in porting codes to the new computer. Too many choices in software or programming environments generally adds confusion for users. Another example of making the process more difficult through additional funding might be the support of a programming environment which is too complex for the occasional user.

The largest impact funding will have on JPL in this endeavor is the lack of funding to hire additional staff for supporting users of the T3D. Most of our funding will be spent in the training of staff so that they can in turn train and assist users with **code** conversion efforts. We will be supported in this effort by having additional Cray applications support personnel at JPL that can assist as we **begin** supporting our users.

What are the Users' Expectations?

What is the main purpose of the project as seen by the user base? Is it geared towards research and development? Production? Do users want to continue to do work exactly as they have in **the** past? Are they willing to change their methods of operation to better utilize the new machine? How much work are they willing to do to take advantage of the new system?

The CRAY T3D will be the production **supercomputer** at JPL. Because the T3D will be a **production** system, users will expect it to be stable, relatively easy to use, and well supported by User Services. Our current system is highly stable, generally going down only for scheduled preventative and software maintenance.

In order to take full advantage of the new T3D, users will need to change the way they operate. They need assistance in the form of classes and documentation to learn how to change from their current modes of operation to ones that will prove more productive on the T3D. How much user effort this will require is unknown at this time, but it will require user effort.

Are Other Facilities Available to the Users?

Do the targeted users of the proposed system have access to other services/facilities for their computing **needs**? If other services exist, what will differentiate this service from the others? Are the potential users in need of a different service? Unless the proper services are offered, the facility will not be used,

JPL users have, and will continue to have, several options to satisfy their **supercomputing** needs. Users requiring vector processing have available through the **Supercomputing** Project limited time at other NASA centers, specifically at the Ames Research Center's Numerical Aerodynamics Simulation Program and at Goddard Space Flight Center's National Center for Computational Sciences. Users requiring parallel computing resources have access to the Intel Touchstone Delta and to an Intel Gamma that serves laboratory personnel. Groups at JPL and the Caltech campus are able to purchase their own computers if they have the required funding. Recently, one group from campus purchased a Kendall Square and a JPL group purchased a 32-node Connection Machine.

As workstation speeds increase relative to **conventional computer/supercomputer** speeds, many users can often satisfy their individual needs through these computers and don't have to use central **supercomputers** to achieve their daily computing goals.

The most compelling reason for users to use a central computer system is for the support services that are provided. These services are extensive and

only begin with user support. In addition, there are system programming, system administration, archiving and operating services that must be provided when a decision is made to break away from the central service. Unless all of these services remain at a high level, users may well abandon the central system for a fast workstation or a group-owned computer over which they maintain control.

JPL as a Center of Excellence

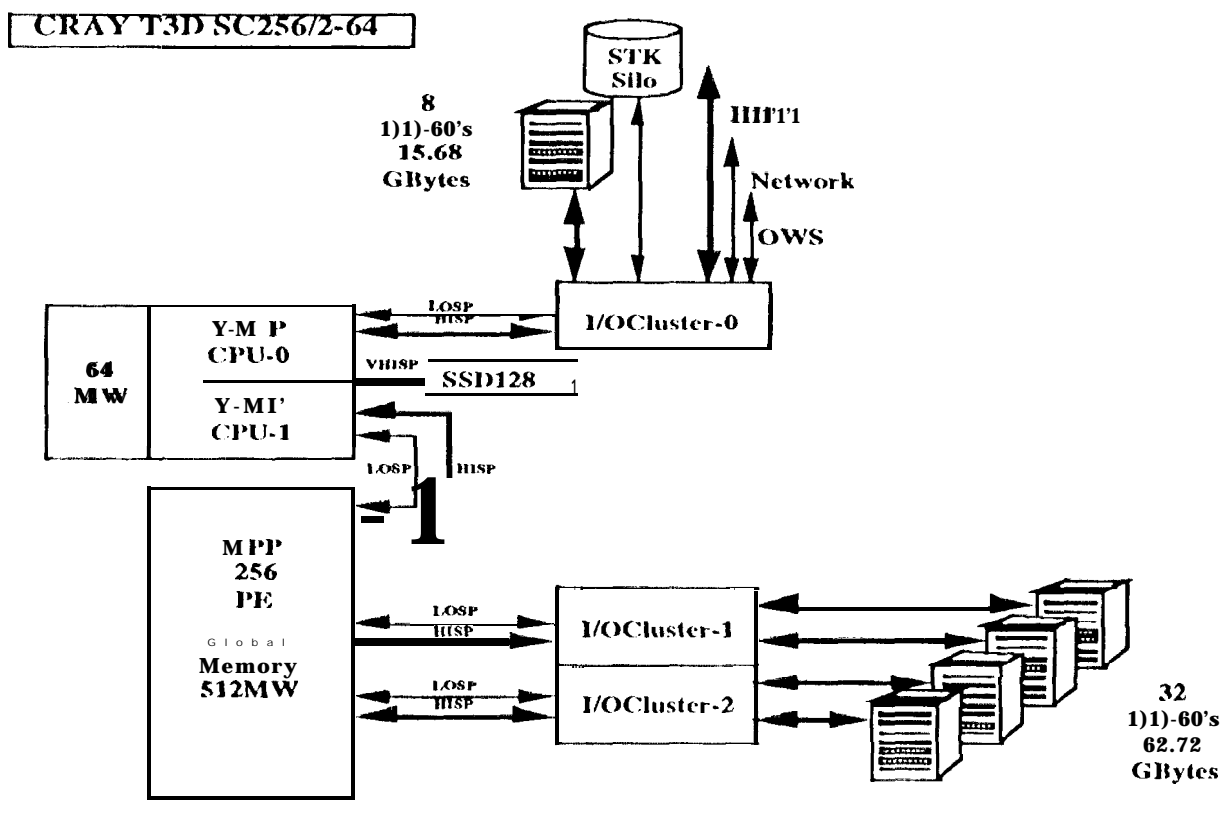
The Hardware

The planning for the T3D and the Center of Excellence is well underway. The production configuration of the T3D in April 1994 is expected to have 256 processor elements with 512

megawords of memory. There will be 62.8 GB of disk space attached directly to the T3D. The T3D will be hosted by either a CRAY M-90 or CRAY SC256 with two processors. An additional 15.7 gigabytes of disk will be attached to the second processor of the host computer and will be used to help satisfy JPL's vector processing needs.

From the arrival of the T3D in October 1993 until April 1994 (longer if necessary), the current CRAY Y-MP2E will be available to handle the production needs of the user community. This overlap will allow time for users to convert codes to execute on the T3D.

A diagram of the April 1994 configuration is shown below.



Where Are We Now?

JPL is just beginning the process. As of the end of February, the staff member with the most MPP experience has attended the CRAY T3D Emulator class. On returning from Eagan, she gave a one hour

overview to the other members of the staff and has also made her training materials available to us. Additional staff will take the T3D Emulator class as slots become available and all staff should be able to attend this training before the T3D is delivered in October. As each person returns from training, he

only begin with user support. In addition, there are system programming, system administration, archiving and operating services that must be provided when a decision is made to break away from the central service. Unless all of these services remain at a high level, users may well abandon the central system for a fast workstation or a group-owned computer over which they maintain control.

JPL as a Center of Excellence

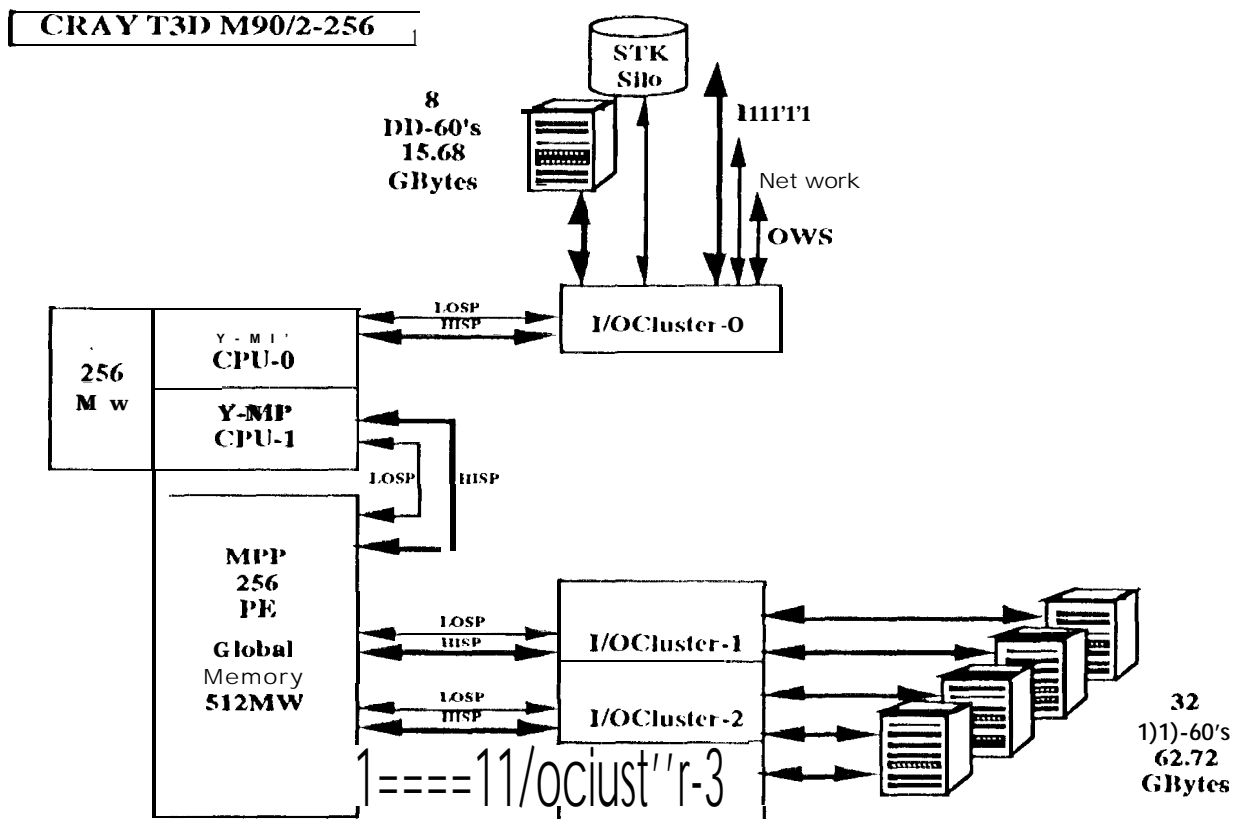
The Hardware

The planning for the T3D and the Center of Excellence is well underway. The production configuration of the T3D in April 1994 is expected to have 256 processor elements with 512

megawords of memory. There will be 62.8 GB of disk space attached directly to the T3D. The T3D will be hosted by either a CRAY M-90 or CRAY SC256 with two processors. An additional 15.7 gigabytes of disk will be attached to the second processor of the host computer and will be used to help satisfy JPL's vector processing needs,

From the arrival of the T3D in October 1993 until April 1994 (longer if necessary), the current CRAY Y-MP2E will be available to handle the production needs of the user community. This overlap will allow time for users to convert codes to execute on the T3D.

A diagram of the April 1994 configuration is shown below.



Where Are We Now?

JPL is just beginning the process. As of the end of February, the staff member with the most MPP experience has attended the CRAY T3D Emulator class. On returning from Eagan, she gave a one hour

overview to the other members of the staff and has also made her training materials available to us. Additional staff will take the T3D Emulator class as slots become available and all staff should be able to attend this training before the T3D is delivered in October. As each person returns from training, he

or she will give an updated presentation to the other members of the staff. This serves several purposes: the staff member solidifies his knowledge by preparing and giving the presentation; the other staff members absorb some new knowledge; everyone keeps up with the changes that are occurring in the product and in the training class; and it helps to form a basis for our user training sessions and materials.

We will install the emulator as soon as it is available following the upgrade of our current system in April 1993. This upgrade will provide us with resources to use the T3D Emulator without impacting current production levels on the Y-MP2E. We will add an additional CPU and 16 MW of memory plus additional

disk with increased I/O bandwidth capacity. The emulator will be used only by User Services staff and by some of the initial users who are part of the Center of Excellence proposal. This should provide the staff a flavor of the MPP and be a tool to help in preparing early documentation and training.

We have identified certain user codes that are good candidates for parallelization where the owners are not experienced users of any MPP. By using the emulator, we hope to get some experience with conversion of real codes once we have experience with the software and concepts and prior to installing the T3D. This experience will feed back into our documentation and training materials.



User Support for an MPP, JPL Begins the Process

Barbara Horner-Miller

California Institute of Technology

Jet Propulsion Laboratory

March 1993

Introduction

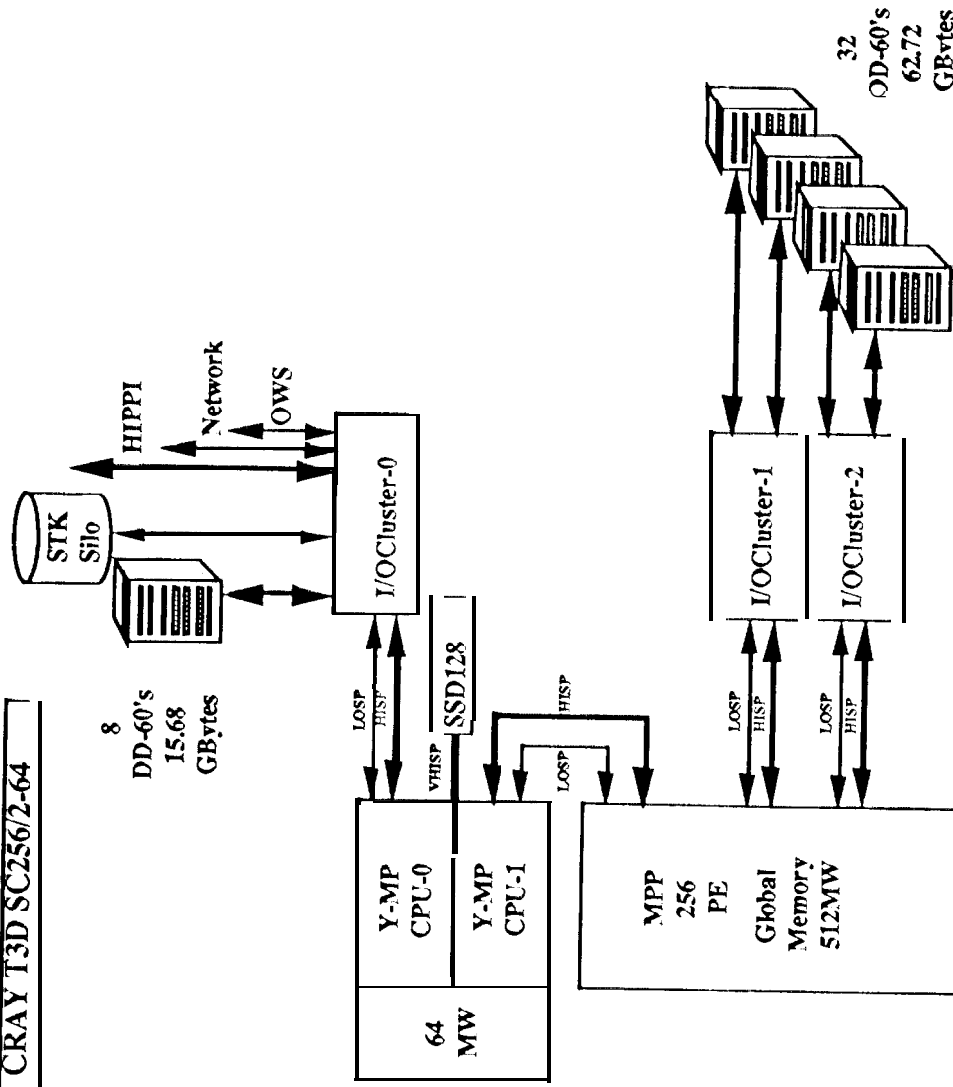
- JPL's Relationship with NASA and the California Institute of Technology**
- Introduction of a Supercomputer, 19S9**
 - Organizational Structure**
 - Services Provided**
 - No User Charges**
- T3D Arrival Slated for October 1993**

Services Currently Provided

- .Consulting Hot Line**
- User Training and Documentation**
- .CRAY System Administration**
- Software Libraries**
- .Visualization Assistance**
- .Code Conversion/Optimization**

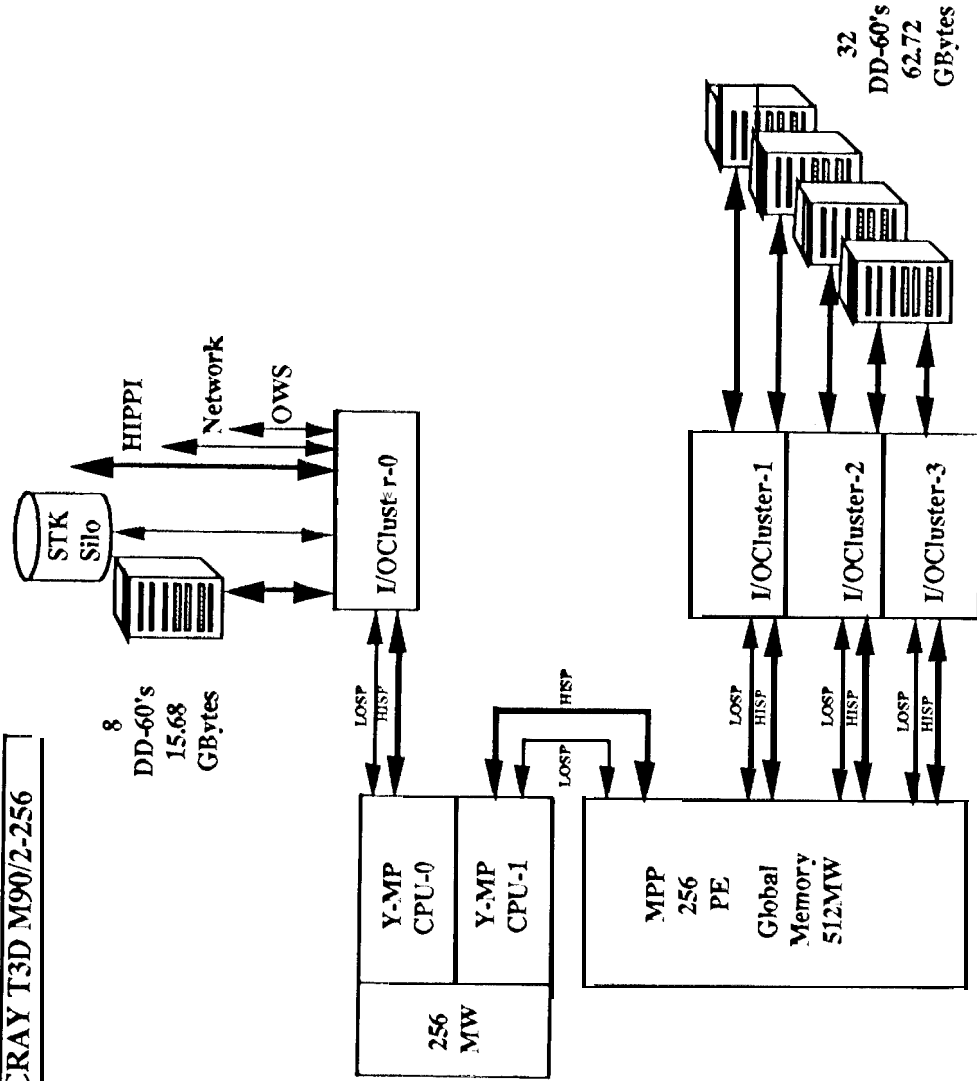
JPL/CALTECH T3D, April 1994

CRAY T3D SC256/2-64



JPL/CALTECH T3D, April 1994

CRAY T3D M90/2-256



Issues Affecting User Support

- .Who Is the Expected User Base?**
- What Is the Available Funding Level?**
- What Are the Users' Expectations?**
- . What Alternatives Are Available to Users?**

Who Is the User Base?

- How Diverse Are the Disciplines
- What Specific Disciplines Are Supported
- What Level of Expertise Do Users Have
 - Computer Programming Expertise
 - MPP Experience

What Available Funding Level?

- Tradeoffs: What You Can Afford vs What You Want
- Level of User Involvement: Total vs Advisory
- Money Can Help or Hinder Your Goals

What Are User Expectations?

- Production vs Research Facility
- Willingness to Change Mode of Operation for better Efficiency
- Two Distinct Phases in Center of Excellence
 - Phase I - Experienced MPP Users
 - Phase II - Novice MPP Users

What Alternatives Are Available?

- What Other Options are Available
- Are Outside Supercomputers Available
- "No Supercomputing" Is an Alternative

Where Are We Now ?

- .Beginning the Process**
- .Emulator and Additional Resources**
- .CRI Courses for All Staff**
- User Support is a Continuum**
 - .Training - Classes vs Individual Assistance**
 - Documentation - Vendor Only vs Supplemental**
 - .Other Services**